

Finalizing Tentative Matches from Truncated Preference Lists

Hisao Tamaki

Department of Computer Science, Meiji University
tamaki@cs.meiji.ac.jp

Abstract. Consider the standard hospitals/residents problem, or the two-sided many-to-one stable matching problem, and assume that the true preference lists of both sides are complete (containing all members of the opposite side) and strict (having no ties). The lists actually submitted, however, are truncated. Let I be such a truncated instance. When we apply the resident-proposing deferred acceptance algorithm of Gale and Shapley to I , the algorithm produces a set of tentative matches (resident-hospital pairs). We say that a tentative match in this set is *finalizable in I* if it is in the resident-optimal stable matching for every completion of I (a complete instance of which I is a truncation). We study the problem we call FTM (Finalizability of Tentative Matches) of deciding if a given tentative match is finalizable in a given truncated instance. We first show that FTM is coNP-complete, even in the stable marriage case where the quota of each hospital is restricted to be 1. We then introduce and study a special case: we say that a truncated instance is *resident-minimal*, if further truncation of the preference lists of the residents inevitably changes the set of tentative matches. Resident-minimal instances are not only practically motivated but also useful in computations for the general case. We give a computationally useful characterization of negative instances of FTM in this special case, which, for instance, can be used to formulate an integer program for FTM. For the stable marriage case, in particular, this characterization yields a polynomial time algorithm to solve FTM for resident-minimal instances. On the other hand, we show that FTM remains coNP-complete for resident-minimal instances, if the maximum quota of the hospitals is 2 or larger. We also give a polynomial-time decidable sufficient condition for a tentative match to be finalizable in the general case. Simulations show that this sufficient condition is extremely useful in a two-round matching procedure based on FTM for a certain type of student-supervisor markets.

1 Introduction

In many matching markets in practice which are modeled by the two-sided, many-to-one stable matching problem of Gale and Shapley [1] (often called the *hospitals/residents problem*), the task of the participants to form a preference list is far from trivial. The large number of participants in the opposite side makes it practically impossible to evaluate all of them in enough details to precisely

determine the preference order. In some cases, quite costly procedures such as interviews are involved in the evaluation process, which makes it even harder to form a precise preference list. If each participant is required to submit a complete preference list in such a market, then the submitted list would be inevitably inaccurate. Otherwise, the lists submitted would be short.

In the latter case, non-negligible number of participants remain unmatched after the matching procedure. This is the case, for example, in the National Resident Matching Program (NRMP) of the United States [4], which assigns candidates for residency to positions in hospitals. To provide further opportunities for the candidates and positions that have failed to be matched in the main matching round called MRM (Main Residency Match), NRMP organizes a post-match program called SOAP (Supplemental Offer and Acceptance Program). Unfortunately, the design of this post-match program lacks a theoretical basis. One clear drawback is that the final matching which results from the entire process, MRM followed by SOAP, is not necessarily stable even under the assumption that, for each agent, the preference list submitted for MRM followed by the list used for SOAP, if any, is a truncation of the true preference list, despite the popularity of NRMP as a working example of the stable matching model. Indeed, this drawback does not depend on how SOAP is administered but is simply due to the possibility that a candidate r unmatched in MRM may have lost the chance of being accepted by some hospital h he would list in SOAP because this hospital h has been filled in MRM by candidates possibly less preferred by h to r .

The algorithmic question studied in this paper is motivated by an approach to address the above issues: a multi-round stable matching procedure. This procedure is designed to produce a stable matching as the final outcome, while allowing participants to incrementally form their preference lists. In the first round, each participant submits a truncation of its true preference list, listing only a small number of candidates it ranks the highest. The deferred-acceptance (DA) algorithm of Gale and Shapley [1] is applied to these truncated preference lists and stops prematurely with a partial outcome. In the second and successive rounds, the lists of the participants are extended, with more agents in the opposite side added in their tails as needed to continue the execution of the DA algorithm.

We list potential benefits of such a multi-round procedure.

1. In the first round, the participant can concentrate on the evaluations of those candidates that are potentially ranked the highest, which makes it easier to form an accurate list for the first submission.
2. In subsequent rounds, some participants do not need to submit extensions to their lists, as those extensions are not required by the DA algorithm. The saving in the evaluation effort for those participants can be huge.
3. Even for a participant who does need to submit an extension, the evaluation task can be easier since (1) some of the remaining candidates may be excluded from considerations as it can be deduced (by the central agency) from the submissions so far that they have no possibility to be matched to the

participant and (2), except for in the last round, the participant may concentrate on those candidate it considers the strongest among the remaining candidates, similarly to the situation in the first round.

We model the situation after each round in such a multi-round procedure as follows. The matching is to be made between the set R of residents and the set H of hospitals. We assume that the true preference lists of both the residents and hospitals are complete, listing all members on the opposite side, and strict, allowing no ties. If J is an instance with complete preference lists and I is obtained from J by truncating some preference lists in J , then we call I a *truncation* of J and J a *completion* of I . We allow truncations on both sides.

When we run the resident-proposing DA algorithm on an instance I executing as many steps as possible in the absence of the missing parts of the preference lists, the execution results in the set of *tentative matches* for I , which we denote by $\mathbf{tent}(I)$. Here, and throughout the paper, a match means simply a resident-hospital pair and should not be confused with a matching, which is a set of matches with certain properties. To deal with the truncations of the preference lists of hospitals to suit our purposes, we need some adaptation of the standard DA algorithm: hospital h may reject the proposal of resident r not in its preference list only when h has filled its quotas by residents which do appear in its list; the proposal of r to h remains *pending* otherwise. See Section 2 for a formal definition of the adapted DA algorithm and tentative matches for truncated instances. If I happens to be complete, then $\mathbf{tent}(I)$ is nothing but the resident-optimal stable matching for I [1]. When I is truncated, each match in $\mathbf{tent}(I)$ is truly tentative and may eventually be rejected when the DA algorithm continues execution on some completion of I . We are interested in the following property of tentative matches and the question on this property.

Definition 1. *We say that a match in $\mathbf{tent}(I)$ is finalizable in I if this match is in $\mathbf{tent}(J)$, the resident-optimal stable matching for J , for every completion J of I .*

FTM (Finalizability of Tentative Matches)

Instance An instance I of the hospitals/residents problem and a match $(r, h) \in \mathbf{tent}(I)$.

Question Is (r, h) finalizable in I ?

For positive integer k , we write k -FTM for the version of FTM where the instances are restricted to those having quota at most k for every hospital: in particular, 1-FTM deals with the one-to-one (stable marriage) instances.

In each round of our multi-round procedure, we first compute $\mathbf{tent}(I)$ by the DA algorithm, where I is the instance specified by the submissions up to that round, and then compute the set of finalizable matches in $\mathbf{tent}(I)$. Those finalizable matches are officially finalized and, in the successive rounds, residents in the finalized matches stop participation and the quotas of the hospitals therein are reduced. In the final round, each remaining participant is asked to submit the complete preference list on the remaining participants in the opposite side.

It is clear from the definition of finalizability that the matching formed by the entire process is identical to the one that would be obtained in one round where the participants submit the complete true preference lists. It is important to note here that we do not need to compute the set of finalizable matches exactly: any subset is sufficient to ensure the correct outcome and a large subset is desirable for having good progresses through rounds. Thus, even though we have negative results on the tractability of FTM as described below, they by no means deny the utility of the notion of finalizability.

Our first result is indeed negative (Theorem 1): FTM, even 1-FTM in fact, is coNP-complete.

We then look at a special case. When the resident-proposing DA algorithm is executed on an instance and resident r is tentatively matched to hospital h in the outcome, the tail part of the preference list of r after h remains “unconsumed” by the algorithm. We say that an instance I is *resident-minimal* if this unconsumed tail is empty for every resident r . Equivalently, I is resident-minimal if the set of matches (r, h) such that h is on the preference list of r equals the set of matches that are proposed in the execution of the DA algorithm on I . Resident-minimal instances are of interest for the following reasons.

1. A natural and purely algorithm-driven matching procedure with incremental submissions would ask for further submissions of participants only when extending their preference lists is absolutely necessary for a progress. In a procedure that applies this policy on residents, the instance we have at each execution step is resident-minimal.
2. Suppose we use a backtrack algorithm to decide if a match is finalizable in a general truncated instance I , which executes the DA algorithm and branches on the next preferred hospital of a resident when it is not given in I . The extension of I that the algorithm constructs in each branching path eventually becomes resident-minimal and the backtrack search beyond this search node can be pruned if an efficient algorithm for resident-minimal instances is available.
3. Each general truncated instance I has a further truncation I' that is resident-minimal. The finalizability in I' is a sufficient condition for the finalizability in I and therefore an efficient computation for resident minimal instances would be useful in estimating the set of finalizable matches in the general case.

Although it is possible to define an analogous notion of hospital-minimal instances, it is not as natural or as useful as that of resident-minimal instances mainly because the preference lists of hospitals are not “sequentially consumed” in the resident-proposing algorithm.

We write FTM-RM for FTM (and k -FTM-RM for k -FTM) in which the instances are restricted to be resident-minimal. Our main result on FTM-RM is a computationally useful characterization of negative instances of FTM-RM (Theorem 2). This characterization may be used, for example, to formulate an integer program or to design a dynamic programming algorithm for FTM-RM.

We also give a polynomial time algorithm for 1-FTM-RM, the stable marriage case, based on this characterization (Theorem 4). On the other hand, we show that 2-FTM-RM, and hence FTM-RM, remains coNP-complete (Theorem 5).

We also develop a polynomial-time decidable sufficient condition for a tentative match being finalizable (Theorem 6). This sufficient condition may be used to compute a subset of finalizable matches in the proposed multi-round stable matching procedures. We also show that this condition is necessary for 1-FTM-RM (Theorem 7). This gives another proof that 1-FTM-RM is polynomial time solvable.

Applications As in typical Japanese Universities, every student in the author’s department is required to take a one-year research project course before graduation, supervised by one of the faculty members. This situation gives rise to a typical instance of a many-to-one stable matching problem.

Since 2014, the department has been administering a two-round stable matching procedure. Supervisors first submit a complete rank list of students based on grade scores and the results of interviews. Then, in the first round of matching, each student lists up to 3 most preferred supervisors in their rank list. The DA algorithm is executed on this truncated instance and the resulting tentative matches are tested for finalizability using the sufficient condition described in Section 5. Only those students without finalized matches proceed to the second round, where they submit a complete preference list of the supervisors who are not filled by finalized matches.

This two-round procedure has been working well: a somewhat surprisingly large number of students are finalized in the first round, resulting in a huge amount of saving in the evaluation effort and stress on the students’ side. Due to the lack of publicly disclosable statistics, we perform simulations in Section 6 to reproduce the phenomenon in a transparent manner.

It would be a challenging research topic to study the feasibility of replacing the current NRMP procedure by a two-round matching procedure in our approach. The advantage of having a stable matching as a final outcome is attractive and the success in the smaller scale market described above is encouraging. The first step of the feasibility study would be to compute the finalizability of the matches produced by the main matching procedure of NRMP in the past, interpreting the preference lists used in the procedure as truncations of the true lists. This task is challenging, because of the intractability of FTM and the size of the market. We note, however, that we may not necessarily need exact solutions. Reasonably good estimates on the number of finalizable matches may be sufficient for our evaluation purposes. Theorem 2 (a characterization of negative instances in the resident-minimal case) and Proposition 8 together with Theorem 6 (a polynomial time computable sufficient condition for finalizability) would be indispensable in computing upper and lower bounds on that number.

Related work Truncation of preference lists in the two-sided matching model have been studied in a different context, namely strategic manipulations. For

example, Roth and Rothblum [7] study the one-to-one matching case and show that there are instances where a participant on the proposed side (a hospital in our model) may expect to benefit significantly by truncating its true preference list, even in the situation where little information on the preference lists of other participants is available.

More traditionally, incomplete preference lists arise not as truncations but as true representations of preferences, where candidates not on the list are simply meant unacceptable. There has been active research on the complexity of computing stable matchings for instances allowing incomplete preference lists and/or ties (see [3] for a survey).

Rastegari *et al.* [6] and Rastegari *et al.* [5] address the incompleteness of the preference orders that are inevitable in large markets in practice. They analyze matching markets where the participants submit their preferences in the form of partial orders that are consistent with their true total orders. The authors of [6] aim at optimizing the number of interviews needed to sufficiently refine the partial orders while the authors of [5] study the complexity of reasoning about the stable matchings for the true hidden preference orders using the partial information available. Our present work may be viewed as dealing with a special case of their model, where the partial preference order can be represented in the form of a truncation. Although successively extending preference lists is restrictive than successive refinements of partial orders, it allows the participants to concentrate on selecting top preferences first and help reduce the chance of regrets in the submitted lists, compared to the case where complete lists are required in one shot. The advantage of being restrictive is that we may have more computationally positive results: none of the positive results in this paper seem to extend to the general partial order model.

A result analogous to the coNP-completeness of 1-FTM (Theorem 1) may be found in their work [5]. Restricting themselves to the stable marriage case, they consider the problem, among others, of deciding if a given match is a *necessary match*, that is, if it is contained in the resident-optimal (employer-optimal, in their setting) stable matching for every completion of the given partial orders into total orders. They show that this problem is coNP-complete. Theorem 1 in our present paper implies that this hardness holds for special instances where the partial orders are restricted to those representable by truncations. Their result does not imply our result and, moreover, neither does their proof, since their reduction uses partial orders that are not representable by truncations.

The rest of this paper is organized as follows. Section 2 gives preliminaries of this paper. In Section 3 we prove the coNP-completeness of 1-FTM. In Section 4 we study resident-minimal instances. In Section 5, we give the sufficient condition for finalizability. In Section 6, we describe our simulation results on the student-supervisor assignment problem. We conclude the paper in Section ??, where we point to some directions for future work.

2 Preliminaries

Formally, an instance I of our hospitals/residents problem is a 5-tuple $(R, H, \{q_h\}_{h \in H}, \{\lambda_r\}_{r \in R}, \{\pi_h\}_{h \in H})$, where R is the set of *residents*, H is the set of *hospitals*, q_h for each h is the quota of h , λ_r for each $r \in R$ is the preference list of r on H , and π_h for each $h \in H$ is the preference list of h on R . The first three components R , H , and $\{q_h\}_{h \in H}$ will always be denoted by these symbols and when we speak of various instances in a context, these components will be common among those instances: only the preference lists will vary. A preference list on a set S is *complete* if it lists all members of S . Our instances in general may have preference lists that are not complete. An instance is *resident-complete* (*hospital-complete*, resp.) if the preference list of each *resident* (*hospital*, resp.) is complete. It is *complete* if it is both resident- and hospital-complete. A *match* is a pair in $R \times H$: we say match (r, h) *involves* r and h . For each set of matches M and $h \in H$, we define $\mathbf{res} M = \{r \mid (r, h) \in M \text{ for some } h \in H\}$ and $\mathbf{res}_h M = \{r \mid (r, h) \in M\}$. A set M of matches is a *matching* in I if M contains at most one match that involves r , for each $r \in R$, and $|\mathbf{res}_h M| \leq q_h$ for each $h \in H$.

We use $+$ operator for concatenation of sequences and for appending or prepending elements to sequences. A sequence α is a *prefix* of a sequence β , and β is an *extension* of α , if β can be written as $\alpha + \gamma$ for some possibly empty sequence γ . The length of sequence α is denoted by $|\alpha|$. Our sequences will never have duplicate elements and therefore the length of α is precisely the cardinality of the set of elements in α .

We say that an instance I is an *extension* of an instance J , and that J is a *truncation* of I , if the preference list of each resident and each hospital in I is the extension of that in J . An extension I of J is *resident-changeless* (*hospital-changeless*, resp.) if the preference list of each resident (hospital, resp.) is identical in I and J .

We apply the DA algorithm to a truncated instance I and stop when further execution is not possible due to the truncations. We represent the intermediate result of this execution by two sets of matches: *tentative matches* are those that have been proposed but not rejected in the algorithm execution; *pending matches* are tentative matches that have not been rejected because of the incompleteness of the preference list of the hospital involved. See below for a more precise definition.

We formalize the execution of our version of the DA algorithm as *event sequences*. Let I be an instance. An *event* for I is either $(r, h)^+$, the *proposal* of a match (r, h) , or $(r, h)^-$, the *rejection* of a match (r, h) . Let σ be a sequence of events for I (or an *event sequence* for I , for short). We say that (r, h) is *proposed* (*rejected*, resp.) in σ if $(r, h)^+$ ($(r, h)^-$, resp.) appears in σ . We denote by $\mathbf{prop}(\sigma)$ ($\mathbf{rej}(\sigma)$, resp.) the set of matches that are proposed (rejected, resp.) in σ . We define $\mathbf{tent}(\sigma) = \mathbf{prop}(\sigma) \setminus \mathbf{rej}(\sigma)$ and call a match in $\mathbf{tent}(\sigma)$ *tentative in* σ . In words, a match is tentative in σ if it is proposed but not rejected in σ . If a tentative match (r, h) in σ is such that r is not in the preference list of h in

the instance I then it is a *pending match* in σ with respect to I . We denote by $\mathbf{pend}_I(\sigma)$ the set of pending matches in σ with respect to I .

Remark 1. In the standard DA algorithm which regards a missing resident r in the preference list of hospital h as unacceptable to h , there is no notion of pending matches: the proposal of r to h can be immediately rejected. In our version, such a proposal should be pending, unless h has filled its quota with proposals from residents in its list, as we intend to continue the algorithm when more residents are added to the preference list of h . Also note that the set $\mathbf{tent}(\sigma)$ depends only on the event sequence σ and not on the instance while the set $\mathbf{pend}_I(\sigma)$ depends on the instance.

Let I be an instance and M an arbitrary set of matches. We say that $(r, h) \in M$ is *ousted from M in I* , if the preference list of h contains at least q_h residents from $\mathbf{res}_h M$ and either r is missing from this list or preceded by q_h or more residents from $\mathbf{res}_h M$ in this list. In other words, (r, h) is ousted from M in I , if, no matter how the preference list of h in I is completed, r is not among the top q_h members of M in the completed list. We let $\mathbf{ousted}_I(M)$ denote the set of matches ousted from M in I . We say that an event sequence is *I -feasible* if it can be shown so by the inductive procedure below.

1. An empty sequence is I -feasible.
2. Suppose an event sequence σ is I -feasible. Then, $\sigma + (r, h)^+$ for each $(r, h) \in R \times H$ is I -feasible if $r \notin \mathbf{res}_{\mathbf{tent}(\sigma)}$, $(r, h) \notin \mathbf{prop}(\sigma)$, h appears in the preference list of $r \in I$, and $(r, h') \in \mathbf{rej}(\sigma)$ for every $h' \in H$ that precedes h in the preference list of r in I . On the other hand, $\sigma + (r, h)^-$ for each $(r, h) \in R \times H$ is I -feasible if (r, h) is in $\mathbf{ousted}_I(\mathbf{prop}(\sigma)) \setminus \mathbf{rej}(\sigma)$.

Remark 2. We have $\mathbf{ousted}(\mathbf{prop}(\sigma)) \setminus \mathbf{rej}(\sigma) = \mathbf{ousted}(\mathbf{tent}(\sigma))$ for I -feasible σ . Therefore, the condition for the I -feasibility of $\sigma + (r, h)^-$ above may be expressed as $(r, h) \in \mathbf{ousted}(\mathbf{tent}(\sigma))$, which is more consistent with the traditional definition of the DA algorithm. We use the condition in the present form, since it makes the monotonicity of the feasibility expressed by the following proposition obvious.

For brevity, we refer to I -feasible event sequences simply as I -feasible sequences.

Proposition 1. *Let I be an instance and let $\sigma + e$ be an I -feasible sequence where e is a single event. Then, for each I -feasible sequence σ' that contains all events in σ but not e , $\sigma' + e$ is I -feasible.*

Given instance I , the execution of the DA algorithm on I results in an arbitrary, due to the non-determinacy of the algorithm, but maximal I -feasible sequence. The following observation that is well known for the standard DA algorithm [1] holds also for our variant.

Proposition 2. *Let I be an instance and let σ and σ' be two maximal I -feasible sequences. Then, σ and σ' contain the same set of events.*

Proof. Suppose σ' contains an event that is not in σ . Among such events, choose one that appears first in σ' and call it e . Then σ followed by e is I -feasible, due to Proposition 1, contradicting the maximality of σ . Therefore, σ' does not have any event not in σ and vice versa. \square

For an instance I with a maximal I -feasible sequence σ , we denote $\mathbf{prop}(\sigma)$, $\mathbf{rej}(\sigma)$, $\mathbf{tent}(\sigma)$, and $\mathbf{pend}_I(\sigma)$ by $\mathbf{prop}(I)$, $\mathbf{rej}(I)$, $\mathbf{tent}(I)$, and $\mathbf{pend}(I)$. This notation is justified since these sets do not depend on the choice of σ and determined solely by I , by Proposition 2. We say that instance I *proposes* (*rejects*, resp.) a match if it is in $\mathbf{prop}(I)$ ($\mathbf{rej}(I)$, resp.).

We say that an event sequence is *feasible* if it is I -feasible for some instance I .

3 Hardness of 1-FTM

To prove hardness results we use the following folklore. A similar statement appears in the description of SATISFIABILITY problem in Garey and Johnson [2]. We include a proof for self-containedness.

Proposition 3. *SAT is NP-complete, even when restricted to a clause set in which each variable appears exactly twice positively and exactly once negatively.*

Proof. Let a clause set S be given. We show below that S can be converted into a clause set S' , without changing the satisfiability, in which each variable appears exactly three times and moreover at least once positively and at least once negatively. By replacing some variables by their negations if necessary, S' may further be converted into a clause set satisfying the condition of the proposition.

Suppose variable x occurs k times in S . We may assume $k \geq 2$ since otherwise the value of x can be fixed without changing the satisfiability. Then, we replace occurrences of x by distinct new variables x_i , $1 \leq i \leq k$, and add clauses $\bar{x}_i \vee x_{i+1}$ for $1 \leq i \leq k$, where $x_{k+1} = x_1$, to force these variables to take the same value. The clause set S' is obtained from S by doing this for all variables. \square

Theorem 1. *1-FTM, and hence FTM, is coNP-complete.*

Proof. That FTM is in coNP is trivial. We prove the hardness by reducing SAT to the complement of 1-FTM.

Let S be an arbitrary set of SAT clauses. Let X be the set of variables of S , and C_1, \dots, C_m the enumeration of clauses in S . Relying on Proposition 3, we assume that each variable occurs positively in exactly two clauses and negatively in exactly one clause.

We construct an instance $I = (R, H, \{q_h\}_{h \in H}, \{\pi_h\}_{h \in H}, \{\lambda_r\}_{r \in R})$ as follows. R consists of two distinguished residents r_0 and r_1 , together with distinct residents $r_x^1, r_x^2, p_x^0, p_x^1$, and p_x^2 for each $x \in X$. Fix $x \in X$. Let C_{j_0}, C_{j_1} , and C_{j_2} be the three clauses in which x appears and assume that the occurrence of x in C_{j_0} is negative. We say that resident p_x^i , $i = 0, 1, 2$, is *associated with* the

occurrence of x in C_{j_i} . H consists of a distinguished hospital h_0 together with distinct hospitals h_x^{-1} , h_x^{-2} , h_x^1 , and h_x^2 for each $x \in X$ and distinct hospitals h_j for each $1 \leq j \leq m$. The preference lists of the hospitals are as follows. The unspecified parts of the lists are immaterial.

1. The list of h_0 starts with r_0 followed by r_1 .
2. For $x \in X$ and $i = 1, 2$, the list of h_x^{-i} starts with r_x^i followed by p_x^0 .
3. For $x \in X$ and $i = 1, 2$, the list of h_x^i starts with r_x^i followed by p_x^i .
4. For $1 \leq j \leq m$, the list of h_j starts with the residents (of the form p_x^0 , p_x^1 , or p_x^2 for some x) that are associated with the variable occurrences in C_j , in an arbitrary order, followed by r_0 .

The preference lists of the residents are as follows. Unlike in the description above for hospitals, these lists are truncated after the specified elements.

1. The list of r_0 is h_1, h_2, \dots, h_m , followed by h_0 .
2. The list of r_1 consists solely of h_0 .
3. For each $x \in X$, the lists of r_x^1 and r_x^2 are empty.
4. For each $x \in X$, the list of p_x^0 starts with h_x^{-1} and h_x^{-2} in this order, followed by h_j , where C_j is the clause that contains the variable occurrence with which resident p_x^0 is associated.
5. For each $x \in X$ and $i = 1, 2$, the list of p_x^i starts with h_x^i followed by h_j , where C_j is the clause that contains the variable occurrence with which resident p_x^i is associated.

See Figure 1 for an example.

First observe that $\text{tent}(I) = \{(r_1, h_0), (r_0, h_1)\} \cup \{(p_x^0, h_x^{-1}), (p_x^1, h_x^1), (p_x^2, h_x^2) \mid x \in X\}$. We show that (r_1, h_0) is not finalizable in I if and only if S is satisfiable.

Let J be an extension of I . We say that resident p of the form p_x^i is *activated* in J , if match (p, h_j) is proposed in J , where h_j is such that C_j contains the variable occurrence to which p is associated and hence h_j is the last entry of the preference list of p in I . Observe that, p_x^0 is activated if and only if r_x^1 chooses h_x^{-1} and r_x^2 chooses h_x^{-2} as their first hospitals on their lists. Similarly, p_x^i , $i = 1, 2$, is activated if and only if r_x^i chooses h_x^i . Therefore, for each $x \in X$, the two events (1) p_x^0 is activated and (2) both p_x^1 and p_x^2 are activated are mutually exclusive and, moreover, we may choose the way the lists of r_x^1 and r_x^2 are extended so that at least one of (1) and (2) happens. Thus, the activation of residents p_x^0 , p_x^1 , and p_x^2 can properly simulate the truth assignment to variable x .

Also observe that (r_1, h_0) is rejected if and only if (r_0, h_0) is proposed, which happens if and only if there is a chain of rejections/proposals of resident r_0 through the hospitals h_1, \dots, h_m leading to this proposal. Since the pair (r_0, h_j) is rejected, provided that this pair is proposed, if and only if at least one resident on the list of h_j that is associated with a variable occurrence in C_j is activated, we conclude that S is satisfiable if and only if there is an extension of I in which (r_1, h_0) is rejected. \square

4 Resident-minimal instances

In this section, we study resident-minimal instances.

			$r_0: h_1 h_2 h_3 h_4 h_0$ $r_1: \mathbf{h_0}$
		$h_x^{-1}: r_x^{-1} p_x^0$ $h_x^{-2}: r_x^{-2} p_x^0$ $h_x^{-1}: r_x^{-1} p_x^1$ $h_x^{-2}: r_x^{-2} p_x^2$	$r_x^1:$ $r_x^2:$ $p_x^0: \mathbf{h_x^{-1} h_x^{-2} h_1}$ $p_x^1: \mathbf{h_x^{-1} h_2}$ $p_x^2: \mathbf{h_x^{-2} h_3}$
$C_1: \sim x \vee y$	$h_0: r_0 r_1$	$h_y^{-1}: r_y^{-1} p_y^0$ $h_y^{-2}: r_y^{-2} p_y^0$ $h_y^{-1}: r_y^{-1} p_y^1$ $h_y^{-2}: r_y^{-2} p_y^2$	$r_y^1:$ $r_y^2:$ $p_y^0: \mathbf{h_y^{-1} h_y^{-2} h_3}$ $p_y^1: \mathbf{h_y^{-1} h_1}$ $p_y^2: \mathbf{h_y^{-2} h_4}$
$C_2: x \vee \sim z$	$h_1: p_x^0 p_y^1 r_0$	$h_z^{-1}: r_z^{-1} p_z^0$ $h_z^{-2}: r_z^{-2} p_z^0$ $h_z^{-1}: r_z^{-1} p_z^1$ $h_z^{-2}: r_z^{-2} p_z^2$	$r_z^1:$ $r_z^2:$ $p_z^0: \mathbf{h_z^{-1} h_z^{-2} h_2}$ $p_z^1: \mathbf{h_z^{-2} h_3}$ $p_z^2: \mathbf{h_z^{-2} h_4}$
$C_3: x \vee \sim y \vee z$	$h_2: p_x^1 p_z^0 r_0$		
$C_4: y \vee z$	$h_3: p_x^2 p_y^0 p_z^1 r_0$ $h_4: p_y^2 p_z^2 r_0$		
(a)		(b)	

(a) Clause set S

(b) The instance of 1-FTM corresponding to S

Tentative matches are shown in bold face from both sides

Fig. 1. Reduction from SAT to the complement of 1-FTM

4.1 Simple extensions and prescriptions

In this subsection, we define the notions of simple extensions and prescriptions, which characterize negative instances of FTM-RM. Hospital-complete instances play an important role here.

Proposition 4. *Let I be a resident-minimal instance and J a resident-changeless and hospital-complete extension of I . Then J is also resident-minimal. We also have $\mathbf{prop}(I) = \mathbf{prop}(J)$ and $\mathbf{tent}(I) \setminus \mathbf{pend}(I) \subseteq \mathbf{tent}(J) \subseteq \mathbf{tent}(I)$.*

Proof. Let L be the set of all matches (r, h) such that h is on the preference list of r in I . Since I is resident-minimal, we have $\mathbf{prop}(I) = L$. As $\mathbf{prop}(I) \subseteq \mathbf{prop}(J) \subseteq L$, we have $\mathbf{prop}(I) = \mathbf{prop}(J)$ and J is resident-minimal. It immediately follows that $\mathbf{tent}(J) \subseteq \mathbf{tent}(I)$. Let (r, h) be a match in $\mathbf{tent}(I) \setminus \mathbf{pend}(I)$. Then, r is on the preference list of h in I and hence extending the preference list of h does not affect the rank of r . Therefore, we have $(r, h) \in \mathbf{tent}(J)$ and hence $\mathbf{tent}(I) \setminus \mathbf{pend}(I) \subseteq \mathbf{tent}(J)$. \square

Let σ be an event sequence. We say that an extension $\sigma + \tau$ of σ is *simple* if $\mathbf{prop}(\tau) \cap \mathbf{rej}(\tau) = \emptyset$ or, in words, τ never rejects a proposal made in itself. We say that an extension J of instance I is *simple* if the maximal I -feasible sequence has a simple maximal J -feasible extension or, equivalently, $(\mathbf{prop}(J) \setminus \mathbf{prop}(I)) \cap (\mathbf{rej}(J) \setminus \mathbf{rej}(I)) = \emptyset$. The goal of this subsection is to show that, for

each resident-minimal instance I , we do not need to search through all extensions of a maximal I -feasible sequence to decide the finalizability of a match in $\mathbf{tent}(I)$: we need only to look at simple extensions. Indeed, it will turn out that we need only to look at hospital-complete simple extensions.

Proposition 5. *Let I be a resident-minimal instance and J a simple extension of I . Then, there is a simple and hospital-complete extension J' of I such that $\mathbf{rej}(J) \subseteq \mathbf{rej}(J')$.*

Proof. Let I and J be as in the lemma. We assume without loss of generality that J is resident-minimal: if not, take an appropriate truncation. Let J_1 be an arbitrary hospital-complete and resident-changeless extension of J . Since J is resident-minimal, so is J_1 by Proposition 4. We also have $\mathbf{rej}(J) \subseteq \mathbf{rej}(J_1)$ and, moreover, $(\mathbf{rej}(J_1) \setminus \mathbf{rej}(J)) \subseteq \mathbf{pend}(J)$, since J_1 is a resident-changeless extension of J . We construct a simple extension of I by truncating preference lists of residents in J_1 . Let $M = (\mathbf{rej}(J_1) \setminus \mathbf{rej}(J)) \setminus \mathbf{tent}(I)$. For each $(r, h) \in M$, h is the last entry of the preference list of r in J and hence in J_1 , since $(r, h) \in \mathbf{pend}(J)$ and J is resident-minimal. Let J' be obtained from J_1 by, for each match $(r, h) \in M$, removing h from the preference list of r . Then, we have $\mathbf{prop}(J') = \mathbf{prop}(J_1) \setminus M$ and $\mathbf{rej}(J') = \mathbf{rej}(J_1) \setminus M$. For each $(r, h) \in M$, h is not on the preference list of r in I , since $(r, h) \notin (\mathbf{tent}(I) \cup \mathbf{rej}(J)) \supseteq (\mathbf{tent}(I) \cup \mathbf{rej}(I)) = \mathbf{prop}(I)$. Therefore, J' is an extension of I . We claim that it is a simple extension of I . To see this, observe that $\mathbf{rej}(J') \setminus \mathbf{rej}(J) \subseteq \mathbf{tent}(I)$ from the construction of J' . Since no match in $(\mathbf{prop}(J') \setminus \mathbf{prop}(I)) \subseteq (\mathbf{prop}(J) \setminus \mathbf{prop}(I))$ can be in $\mathbf{rej}(J)$ as J is a simple extension of I , no such match can be in $\mathbf{rej}(J')$. Therefore, J' is a simple extension of I . As $M \subseteq (\mathbf{rej}(J_1) \setminus \mathbf{rej}(J))$ and $\mathbf{rej}(J') = \mathbf{rej}(J_1) \setminus M$, we have $\mathbf{rej}(J) \subseteq \mathbf{rej}(J')$ and are done. \square

For the time being, we concentrate on resident-minimal instances that are also hospital-complete and try to characterize their simple extensions.

Proposition 6. *Let I be a resident-minimal and hospital-complete instance and J a simple extension of I . Let $P = \mathbf{prop}(J) \setminus \mathbf{prop}(I)$ and $X = \mathbf{rej}(J) \setminus \mathbf{rej}(I)$. Then, these sets of matches satisfy the following conditions.*

- P1:** $P \cap \mathbf{prop}(I) = \emptyset$.
- P2:** For each $r \in R$, there is at most one $h \in H$ such that $(r, h) \in P$.
- P3:** $X \subseteq \mathbf{tent}(I)$.
- P4:** $\mathbf{res} P \cap \mathbf{res} \mathbf{tent}(I) \subseteq \mathbf{res} X$.
- P5:** For each $h \in H$, we have $|\mathbf{res}_h(P \cup (\mathbf{tent}(I) \setminus X))| \leq q_h$. Moreover, if $\mathbf{res}_h X$ is non-empty then we have $|\mathbf{res}_h(P \cup (\mathbf{tent}(I) \setminus X))| = q_h$.
- P6:** For each $h \in H$, each member of $\mathbf{res}_h(P \cup (\mathbf{tent}(I) \setminus X))$ precedes all members of $\mathbf{res}_h X$ in the preference list of h in I .

\square

A *prescription* for resident-minimal and hospital complete instance I is a pair (P, X) of sets of matches that satisfies the conditions P1 through P6 in Proposition 6. The *target set* of prescription (P, X) , denoted by $\mathbf{tgs}(P, X)$ is

defined by $\mathbf{tgs}(P, X) = \{(r, h) \in X \mid r \notin \mathbf{res} P\}$. The crucial part of the proof of the result in this section is in showing that a certain form of converse of Proposition 6 holds: a prescription (P, X) for I implies a simple extension of I that rejects matches in $\mathbf{tgs}(P, X)$.

Proposition 7. *Let I be a resident-minimal and hospital-complete instance and let (P, X) be a prescription for I . Then, we have $|\mathbf{res} P \setminus \mathbf{res} \mathbf{tent}(I)| \geq |\mathbf{tgs}(P, X)|$.*

Proof. Let $Q = \{(r, h) \in P \mid r \notin \mathbf{res} \mathbf{tent}(I)\}$. By condition P2 for (P, X) being a prescription for I , we have $|Q| = |\mathbf{res} P \setminus \mathbf{res} \mathbf{tent}(I)|$. Our goal is to show that $|Q| \geq |\mathbf{tgs}(P, X)|$.

By condition P5 for (P, X) being a prescription for I , we have $|\mathbf{res}_h X| \leq |\mathbf{res}_h P|$ for each $h \in H$. Therefore, we have $|X| \leq |P|$. On the other hand, let (r, h) be a match in $(r, h) \in P \setminus Q$. Because of condition P4, there is some h' such that $(r, h') \in X$. However, by definition, $(r, h') \notin \mathbf{tgs}(P, X)$ as $r \in \mathbf{res} P$. Therefore, we have $|P \setminus Q| \leq |X \setminus \mathbf{tgs}(P, X)|$. As Q is a subset of P and $\mathbf{tgs}(P, X)$ is a subset of X , we have $|P| - |Q| \leq |X| - |\mathbf{tgs}(P, X)|$. Combining this with $|X| \leq |P|$, we conclude that $|Q| \geq |\mathbf{tgs}(P, X)|$. \square

Lemma 1. *Let I be a resident-minimal and hospital-complete instance and suppose there is a prescription (P, X) for I . Then there is some simple extension J of I such that $\mathbf{prop}(J) \setminus \mathbf{prop}(I) \subseteq P$, $\mathbf{rej}(J) \setminus \mathbf{rej}(I) \subseteq X$, and $\mathbf{tgs}(P, X) \subseteq \mathbf{rej}(J)$.*

Proof. Let I and (P, X) be as in the lemma and σ a maximal I -feasible sequence.

We prove the statement of the lemma by induction on $|P|$. We take $\mathbf{tgs}(P, X) = \emptyset$ as the base case, which includes the case $P = X = \emptyset$. The statement is satisfied with $J = I$ in this case.

For the induction step, suppose $\mathbf{tgs}(P, X)$ is non-empty and let $Q = \{(r, h) \in P \mid r \notin \mathbf{res} \mathbf{tent}(I)\}$. Since $\mathbf{tgs}(P, X)$ is non-empty, Q is non-empty by Proposition 7. Let $\sigma + \tau_1$ be an extension of σ such that τ_1 first lists the proposals of matches in Q in an arbitrary order and then lists all rejections, in an arbitrary order, that are made possible by these proposals (without further chain of proposals and rejections). Then, $\sigma + \tau_1$ is maximal I_1 -feasible where I_1 is the extension of I obtained by appending h in the preference list of r for each $(r, h) \in Q$. We claim that $\mathbf{rej}(\tau_1) \subseteq X$. To see this, fix $h \in H$. By condition P5 for (P, X) being a prescription for I , we have $|\mathbf{res}_h(P \cup (\mathbf{tent}(I) \setminus X))| \leq q_h$ and hence $|\mathbf{res}_h(Q \cup (\mathbf{tent}(I) \setminus X))| \leq q_h$. Thus, no match (r, h) is rejected by τ_1 unless $(r, h) \in X$. If $\mathbf{tgs}(P, X) \subseteq \mathbf{rej}(I_1)$ then we are done with $J = I_1$.

So suppose otherwise, that $\mathbf{tgs}(P, X)$ is not contained in $\mathbf{rej}(I_1)$. Consider the prescription (P_1, X_1) for I_1 where $P_1 = P \setminus Q$ and $X_1 = X \setminus \mathbf{rej}(\tau_1)$. We confirm that this pair is indeed a prescription for I_1 . From condition P1 for (P, X) being a prescription for I , we have $P \cap \mathbf{prop}(I) = \emptyset$. Since $\mathbf{prop}(I_1) = \mathbf{prop}(I) \cup Q$ and $P_1 = P \setminus Q$, it follows that $P_1 \cap \mathbf{prop}(I_1) = \emptyset$, condition P1 for (P_1, X_1) being a prescription for I_1 . Condition P2 immediately follows from the corresponding condition for (P, X) . Since $X \subseteq \mathbf{tent}(I)$ (condition P3 for (P, X)) and $X_1 = X \setminus \mathbf{rej}(\tau_1)$, we have $X_1 \subseteq \mathbf{tent}(I) \setminus \mathbf{rej}(\tau_1) \subseteq \mathbf{tent}(I_1)$, condition P3.

For condition P4, we use the facts that P_1 and Q partition P and that Q and $\mathbf{tent}(I) \setminus \mathbf{rej}(\tau_1)$ partition $\mathbf{tent}(I_1)$. Also using condition P4 for (P, X) that $\mathbf{res} P \cap \mathbf{res} \mathbf{tent}(I) \subseteq X$, we have

$$\begin{aligned}
\mathbf{res} P_1 \cap \mathbf{res} \mathbf{tent}(I_1) &= (\mathbf{res} P \setminus \mathbf{res} Q) \cap \mathbf{res}(Q \cup (\mathbf{tent}(I) \setminus \mathbf{rej}(\tau_1))) \\
&= (\mathbf{res} P \setminus \mathbf{res} Q) \cap (\mathbf{res} Q \cup (\mathbf{res} \mathbf{tent}(I) \setminus \mathbf{res} \mathbf{rej}(\tau_1))) \\
&= \mathbf{res} P \cap (\mathbf{res} \mathbf{tent}(I) \setminus \mathbf{res} \mathbf{rej}(\tau_1)) \\
&= (\mathbf{res} P \cap \mathbf{res} \mathbf{tent}(I)) \setminus \mathbf{res} \mathbf{rej}(\tau_1) \\
&\subseteq \mathbf{res} X \setminus \mathbf{res} \mathbf{rej}(\tau_1) \\
&\subseteq \mathbf{res}(X \setminus \mathbf{rej}(\tau_1)) \\
&= \mathbf{res} X_1.
\end{aligned}$$

Therefore, condition P4 holds.

For conditions P5 and P6, observe that

$$\begin{aligned}
P_1 \cup (\mathbf{tent}(I_1) \setminus X_1) &= P_1 \cup ((Q \cup (\mathbf{tent}(I) \setminus \mathbf{rej}(\tau_1))) \setminus X_1) \\
&= P_1 \cup (((Q \cup \mathbf{tent}(I)) \setminus \mathbf{rej}(\tau_1)) \setminus X_1) \\
&= P_1 \cup (Q \cup (\mathbf{tent}(I) \setminus X)) \\
&= P \cup (\mathbf{tent}(I) \setminus X),
\end{aligned}$$

where we have repeatedly used the disjointness between subsets of P and subsets of X . Therefore, for each $h \in H$, we have $\mathbf{res}(P_1 \cup (\mathbf{tent}(I_1) \setminus X_1)) = \mathbf{res}(P \cup (\mathbf{tent}(I) \setminus X))$ and hence condition P5 for (P_1, X_1) follows from that for (P, X) . Moreover, by condition P6 for (P, X) , each member of $\mathbf{res}_h(P \cup (\mathbf{tent}(I) \setminus X))$ precedes all members of X in the preference list of h in I . Since $X_1 \subseteq X$ and I_1 is an extension of I , it follows that each member of $\mathbf{res}_h(P_1 \cup (\mathbf{tent}(I_1) \setminus X_1))$ precedes all members of X_1 in the preference list of h in I_1 : condition P6 holds. We have confirmed that (P_1, X_1) is indeed a prescription for I_1 .

We note that $\mathbf{tgs}(P_1, X_1) = \mathbf{tgs}(P, X) \setminus \mathbf{rej}(\tau_1)$ is non-empty under our current assumption. Therefore, we may apply the induction hypothesis to instance I_1 and prescription (P_1, X_1) for I_1 to obtain a simple and hospital-complete extension I'_1 of I_1 such that $\mathbf{prop}(I'_1) \setminus \mathbf{prop}(I_1) \subseteq P_1$, $\mathbf{rej}(I'_1) \setminus \mathbf{rej}(I_1) \subseteq X_1$, and $\mathbf{tgs}(P_1, X_1) \subseteq \mathbf{rej}(I'_1)$. We have

$$\begin{aligned}
\mathbf{prop}(I'_1) \setminus \mathbf{prop}(I) &= (\mathbf{prop}(I'_1) \setminus \mathbf{prop}(I_1)) \cup (\mathbf{prop}(I_1) \setminus \mathbf{prop}(I)) \\
&\subseteq P_1 \cup Q \\
&= P,
\end{aligned}$$

$$\begin{aligned}
\mathbf{rej}(I'_1) \setminus \mathbf{rej}(I) &\subseteq (\mathbf{rej}(I'_1) \setminus \mathbf{rej}(I_1)) \cup (\mathbf{rej}(I_1) \setminus \mathbf{rej}(I)) \\
&\subseteq X_1 \cup \mathbf{rej}(\tau_1) \\
&= X,
\end{aligned}$$

and

$$\begin{aligned}\mathbf{tgs}(P, X) &\subseteq \mathbf{tgs}(P_1, X_1) \cup \mathbf{rej}(\tau_1) \\ &\subseteq \mathbf{rej}(I'_1) \cup \mathbf{rej}(\tau_1) \\ &= \mathbf{rej}(I'_1),\end{aligned}$$

since $\mathbf{rej}(\tau_1) \subseteq \mathbf{rej}(I_1) \subseteq \mathbf{rej}(I'_1)$. Therefore, setting $J = I'_1$, the statement of the lemma holds. This completes the induction step and hence the proof of the lemma. \square

Lemma 2. *Let I be a resident-minimal and hospital-complete instance and (r_0, h_0) a match in $\mathbf{tent}(I)$ that is not finalizable in I . Let σ be a maximal I -feasible sequence and $\sigma + \tau$ a shortest feasible extension of σ that rejects (r_0, h_0) . Then, $(\mathbf{prop}(\tau), \mathbf{rej}(\tau))$ is a prescription for I with $\mathbf{tgs}(\mathbf{prop}(\tau), \mathbf{rej}(\tau)) = \{(r_0, h_0)\}$.*

Proof. We set $P = \mathbf{prop}(\tau) \setminus \mathbf{rej}(\tau)$ and $X = \mathbf{rej}(\tau) \setminus \mathbf{prop}(\tau)$. It will turn out that $\mathbf{prop}(\tau) \cap \mathbf{rej}(\tau) = \emptyset$ and hence $P = \mathbf{prop}(\tau)$ and $X = \mathbf{rej}(\tau)$.

We first confirm that (P, X) is a prescription for I . Since $\sigma + \tau$ is feasible, $P \subseteq \mathbf{prop}(\tau)$ is disjoint from $\mathbf{prop}(\sigma) = \mathbf{prop}(I)$: condition P1 holds. Since $P \subseteq \mathbf{tent}(\sigma + \tau)$, for each $r \in R$, there is at most one h such that $(r, h) \in P$: condition P2 holds. Since each match rejected by τ but not already in $\mathbf{tent}(\sigma)$ must be in $\mathbf{prop}(\tau)$, we have $X \subseteq \mathbf{tent}(I)$: condition P3 holds. For condition P4, let $r \in \mathbf{res} P \cap \mathbf{res} \mathbf{tent}(I)$. As (r, h) for some h is proposed in τ , some match $(r, h') \in \mathbf{tent}(I)$ must be rejected in τ and hence in $\mathbf{rej}(\tau) \setminus \mathbf{prop}(\tau) = X$. Therefore, we have $\mathbf{res} P \cap \mathbf{res} \mathbf{tent}(I) \subseteq \mathbf{res} X$.

For conditions P5 and P6, fix $h \in H$. Since $P \cup (\mathbf{tent}(I) \setminus X) = \mathbf{tent}(\sigma + \tau)$, we have $|\mathbf{res}_h(P \cup (\mathbf{tent}(I) \setminus X))| \leq q_h$. Moreover, if $\mathbf{res}_h X$ is non-empty, then τ rejects a match involving h and therefore this inequality is tight. Therefore, condition P5 holds. As each member of $\mathbf{res}_h(\mathbf{tent}(\sigma + \tau))$ precedes all members of $\mathbf{res}_h(\mathbf{rej}(\tau))$, condition P6 holds.

We have $(r_0, h_0) \in \mathbf{rej}(\tau)$ and, from the assumption that τ is chosen to be the shortest, r_0 is not involved in any proposal in τ . Therefore $(r_0, h_0) \in \mathbf{tgs}(P, X)$. A match (r, h) in $\mathbf{tgs}(P, X)$ distinct from (r_0, h_0) would also contradict that assumption, since the rejection of such (r, h) may be removed from τ without affecting the feasibility as r is not involved in any proposal in τ . We conclude that $\mathbf{tgs}(P, X) = \{(r_0, h_0)\}$.

By Lemma 1, there is a simple extension J of I such that $\mathbf{prop}(J) \setminus \mathbf{prop}(I) \subseteq P$, $\mathbf{rej}(J) \setminus \mathbf{rej}(I) \subseteq X$, and $(r_0, h_0) \subseteq \mathbf{rej}(J)$. Let $\sigma + \tau'$ be a maximum J -feasible extension of σ . Then, $\mathbf{prop}(\tau') = \mathbf{prop}(J) \setminus \mathbf{prop}(I) \subseteq P \subseteq \mathbf{prop}(\tau)$ and $\mathbf{rej}(\tau') = \mathbf{rej}(J) \setminus \mathbf{rej}(I) \subseteq X \subseteq \mathbf{rej}(\tau)$. All of these inclusions must in fact be equalities, since otherwise $\sigma + \tau'$ is a feasible extension of σ rejecting (r_0, h_0) that is shorter than $\sigma + \tau$, a contradiction. Therefore, we have $\mathbf{prop}(\tau') = \mathbf{prop}(\tau) = P$ and $\mathbf{rej}(\tau') = \mathbf{rej}(\tau) = X$, finishing the proof of the lemma. \square

We have focused on those resident-minimal instances that are also hospital-complete. The following theorem, however, is on general resident-minimal instances.

Theorem 2. . Let I be a resident-minimal instance and (r_0, h_0) a match in $\text{tent}(I)$. Then, the following three conditions are equivalent.

- (1) Match (r_0, h_0) is not finalizable in I .
- (2) There is a resident-changeless and hospital-complete extension I' of I such that there is a prescription (P, Y) for I' with $(r_0, h_0) \in \text{tgs}(P, Y)$.
- (3) There is a simple extension of I that rejects (r_0, h_0) .

Proof. (3) \Rightarrow (1) is trivial. We show (1) \Rightarrow (2) \Rightarrow (3) below.

(1) \Rightarrow (2): Suppose (r_0, h_0) is not finalizable in I . Let J be an extension of I that rejects (r_0, h_0) . Let J' be an arbitrary resident-changeless and hospital-complete extension of J . We let I' be the resident-changeless and hospital-complete extension of I in which the preference list of each hospital is identical to that in J' . By Proposition 4, I' is resident-minimal. Since J' is an extension of I' and rejects (r_0, h_0) , by Lemma 2, there is a prescription (P, Y) for I' such that $(r_0, h_0) \in \text{tgs}(P, Y)$.

(2) \Rightarrow (3): Let I' and (P, Y) be as in condition (2). By Lemma 1, there is a simple extension J of I' such that $\text{tgs}(P, Y) \subseteq \text{rej}(J)$. Since J is a simple extension of I , we are done. \square

This theorem shows that, for resident-minimal instance I , a triple (P, Y, I') , where I' is a resident-changeless and hospital-complete extension of I and (P, Y) is a prescription for I' , is a certificate that each match in $\text{tgs}(P, Y)$ is not finalizable in I . We seek a more concise certificate and generalize the notion of prescription to general resident-minimal instances.

Let I be a resident-minimal instance. A *prescription for I* is a pair (P, X) of sets of matches that satisfies conditions P1, P2, P3, P4, P5 in Proposition 6 together with the following condition that replaces P6.

P6': For each $h \in H$, the following holds. Each member of $\text{res}_h(P \cup ((\text{tent}(I) \setminus \text{pend}(I)) \setminus X))$ precedes all members of $\text{res}_h(X \setminus \text{pend}(I))$ in the preference list of h in I . Moreover, if $\text{res}_h(X \setminus \text{pend}(I))$ is non-empty then $\text{res}_h \text{pend}(I) \subseteq \text{res}_h X$.

The target set $\text{tgs}(P, X)$ of prescription (P, X) is defined in the same manner as in the special case before: $\text{tgs}(P, X) = \{(r, h) \in X \mid r \notin \text{res } P\}$.

Note that if I is hospital-complete then $\text{pend}(I)$ is empty and hence condition P6' is equivalent to condition P6.

Lemma 3. Let I be a resident-minimal instance, I' a resident-changeless and hospital-complete extension of I , and (P, Y) a prescription for I' . Then, (P, X) , where $X = Y \cup (\text{rej}(I') \setminus \text{rej}(I))$ is a prescription for I .

Proof. Conditions P1 and P2 do not depend on X and therefore follow from those conditions for prescription (P, Y) . Since $Y \subseteq \text{tent}(I') \subseteq \text{tent}(I)$ and $\text{rej}(I') \setminus \text{rej}(I) \subseteq \text{tent}(I)$, condition P3 that $X \subseteq \text{tent}(I)$ holds. For condition P4, let $r \in \text{res } P \cap \text{res } \text{tent}(I)$. If $r \in \text{res } P \cap \text{res } \text{tent}(I')$ then $r \in Y$ by condition P4 for prescription (P, Y) . Otherwise, $r \in \text{res}(\text{rej}(I') \setminus \text{rej}(I)) \subseteq \text{res } X$. Therefore,

condition P4 holds. Condition P5 is equivalent to condition P5 for prescription (P, Y) , since $P \cup (\text{tent}(I) \setminus X) = P \cup (\text{tent}(I') \setminus Y)$. For condition P6, fix $h \in H$. In the preference list of h in I' , each member of $\text{res}_h(P \cup (\text{tent}(I) \setminus X))$ precedes all members of $\text{res}_h Y$ by condition P6 for prescription (P, Y) , and obviously precedes all members of $\text{res}_h(\text{rej}(I') \setminus \text{rej}(I))$. Therefore, it precedes all members of $\text{res}_h X$. If $r \notin \text{res pend}(I)$ then r is already in the preference list of h in I . Therefore, each member of $\text{res}_h(P \cup ((\text{tent}(I) \setminus \text{pend}(I)) \setminus X))$ precedes all members of $\text{res}_h(X \setminus \text{pend}(I))$ in the preference list of h in I . Moreover, suppose some $r \in \text{res}_h(X \setminus \text{pend}(I))$ and some $r' \in \text{res}_h(\text{pend}(I) \setminus X)$. Then, since r' is in $\text{res}_h(P \cup (\text{tent}(I) \setminus X))$ and $r \in \text{res}_h Y$, r' must precede r in the preference list of h in I' . But this is impossible since r is on the preference list of h in I while r' is not, a contradiction. Therefore, if $\text{res}_h(X \setminus \text{pend}(I))$ is non-empty then $\text{res}_h \text{pend}(I) \subseteq \text{res}_h X$: condition P6 holds. \square

Lemma 4. *Let I be a resident-minimal instance and (P, X) is a prescription for I . Then, there is some resident-changeless and hospital-complete extension I' of I such that (P, Y) , where $Y = X \setminus \text{rej}(I')$, is a prescription for I' .*

Proof. For each $h \in H$, arbitrarily complete the preference list of h in I so that the residents in $\text{res}_h(\text{pend}(I) \cap X)$ get the lowest ranks in the completed list. Let the resulting instance be I' . We confirm that (P, Y) is a prescription for I' . Conditions P1 and P2 do not depend on P and therefore follow from those conditions for prescription (P, X) . Since $X \subseteq \text{tent}(I)$ and $\text{tent}(I') = \text{tent}(I) \setminus \text{rej}(I')$, condition P3 that $Y \subseteq \text{tent}(I')$ holds. For condition P4, let $r \in \text{res } P \cap \text{res tent}(I')$. Since $r \in \text{res } P \cap \text{tent}(I)$, we have $r \in \text{res } X$ by condition P4 for prescription (P, X) . Therefore, we have $r \in \text{res } X \cap \text{res tent}(I') = \text{res}(X \setminus \text{rej}(I')) = \text{res } Y$, condition P4. Condition P5 is equivalent to condition P5 for prescription (P, X) , since $P \cup (\text{tent}(I) \setminus X) = P \cup (\text{tent}(I') \setminus Y)$.

To show that P6 holds, let $r \in \text{res}_h(P \cup (\text{tent}(I') \setminus Y)) = \text{res}_h(P \cup (\text{tent}(I) \setminus X))$. Suppose first that $r \notin \text{res}_h \text{pend}(I)$. Then, by condition P6' for (P, X) , r precedes all members of $\text{res}_h(Y \setminus \text{pend}(I)) \subseteq \text{res}_h(X \setminus \text{pend}(I))$ in the preference list of h in I and hence in I' as well. Since r precedes all members of $\text{res}_h \text{pend}(I)$ in the preference list of h in I' by the way I' completes the preference list of h , we conclude that r precedes all members of $\text{res}_h X$ in that preference list. Suppose next that $r \in \text{res}_h \text{pend}(I)$. Then, since $\text{res}_h \text{pend}(I)$, having r as a member, is not contained in $\text{res}_h X$, $\text{res}_h(X \setminus \text{pend}(I))$ is empty, by condition P6' for (P, X) . Therefore r precedes all members in $\text{res}_h X$ in the preference list of h in I' , as those members are placed in the lowest positions. In either case, r precedes all members of $\text{res}_h Y \subseteq \text{res}_h X$ in the preference list of h in I' , that is, condition P6 holds for (P, Y) . \square

Thus, a prescription for a general resident-minimal instance is indeed a certificate for the negative answer to the finalizability of a tentative match.

Theorem 3. *Let I be a resident-minimal instance and (r_0, h_0) a match in $\text{tent}(I) \setminus \text{pend}(I)$. Then, there is a prescription (P, X) for I with $(r_0, h_0) \in \text{tgs}(P, X)$ if and only if there is some resident-changeless and hospital-complete extension I' of I and a prescription (P, Y) for I' with $(r_0, h_0) \in (P, Y)$.*

Proof. Suppose first that there is a prescription (P, X) for I with $(r_0, h_0) \in \mathbf{tgs}(P, X)$. By Lemma 4, there is a resident-changeless and hospital-complete extension I' of I and a prescription (P, Y) for I' such that $Y = X \setminus \mathbf{rej}(I')$. As $X \cap \mathbf{rej}(I') \subseteq \mathbf{pend}(I)$ and $(r_0, h_0) \notin \mathbf{pend}(I)$, $(r_0, h_0) \in \mathbf{tgs}(P, X)$ implies $(r_0, h_0) \in \mathbf{tgs}(P, Y)$. For the converse, suppose that there is a resident-changeless and hospital-complete extension I' of I and a prescription (P, Y) for I' with $(r_0, h_0) \in \mathbf{tgs}(P, Y)$. By Lemma 3, (P, X) , where $X = Y \cup (\mathbf{rej}(I') \setminus \mathbf{rej}(I))$, is a prescription for I . Since $\mathbf{tgs}(P, Y) \subseteq \mathbf{tgs}(P, X)$, we have $(r_0, h_0) \in \mathbf{tgs}(P, X)$. \square

We close this subsection by sketching an integer program (IP) for computing a prescription for a given resident-minimal instance I and a match $(r_0, h_0) \in \mathbf{tent}(I)$. More precisely, the IP captures a triple (P, X, Z) , where (P, X) is a prescription for I with $(r_0, h_0) \in \mathbf{tgs}(P, X)$ and Z is a subset of $\mathbf{pend}(I)$ such that there is a resident-changeless and hospital complete extension J of I with $Z = \mathbf{rej}(J) \setminus \mathbf{rej}(I)$ and $(P, X \setminus Z)$ being a prescription for J .

We only describe the variables in the IP and their intended interpretations. The linear constraints are straightforward to write down based on those interpretations. All the variables are binary. For each match $(r, h) \notin \mathbf{prop}(I)$, we have a variable $p_{r,h}$: $p_{r,h} = 1$ if and only if $(r, h) \in P$. For each match $(r, h) \in \mathbf{tent}(I)$, we have a variable $x_{r,h}$: $x_{r,h} = 1$ if and only if $(r, h) \in X$. For each $h \in H$ and a subset S of $\mathbf{res}_h \mathbf{pend}(I)$, we have a variable $z_{h,S}$: $z_{h,S} = 1$ if and only if $\mathbf{res}_h Z = S$. The objective function is the sum of $p_{r,h}$ over all $(r, h) \in (R \times H) \setminus \mathbf{prop}(I)$, which is minimized. The optimal solution of this IP corresponds to a desired prescription with the smallest cardinality of P .

4.2 Polynomial time algorithm for the stable marriage case

In this subsection, we show that 1-FTM-RM, the finalizability of a tentative match for resident-minimal stable marriage instances, is polynomial time solvable.

Let I be a resident-minimal stable marriage instance. We define a bipartite digraph G_I on vertex sets $T = \mathbf{tent}(I)$ and $P = (R \times H) \setminus \mathbf{prop}(I)$ as follows. Let $(r, h) \in T$ and $(r', h') \in P$. There is an edge from (r, h) to (r', h') if and only if $r = r'$. There is an edge from (r', h') to (r, h) if and only if $h = h'$, both r and r' are on the preference list of h in I , and r' precedes r in that list.

Lemma 5. *Let I be a resident-minimal stable marriage instance and (r_0, h_0) a match in $\mathbf{tent}(I)$. Then, there is a simple extension of I that rejects (r_0, h_0) if and only if there is a directed path in G_I from some root (a vertex without incoming edges) of G_I to (r_0, h_0) .*

Proof. Suppose first that I has a simple extension I' that rejects (r_0, h_0) . Let σ be a maximal I -feasible sequence and $\sigma + \tau$ a maximal I' -feasible sequence. We determine a sequence of matches (r_i, h_i) , $i = 0, 1, \dots$, so that the reversed sequence (r_j, h_j) , $j = i, i-1, \dots, 0$, forms a directed path from (r_i, h_i) to (r_0, h_0)

in I , for each i . We maintain the invariant that if $(r_i, h_i) \in T$ then $(r_i, h_i) \in \mathbf{rej}(\tau)$ and if $(r_i, h_i) \in P$ then $(r_i, h_i) \in \mathbf{prop}(\tau)$. We start with the given match (r_0, h_0) .

Suppose $i \geq 0$ and match (r_i, h_i) has been determined. If (r_i, h_i) is a root of G_I then we are done as we have a desired path from (r_i, h_i) to (r_0, h_0) . Suppose otherwise. First suppose that $(r_i, h_i) \in T$. If $(r_i, h_i) \in \mathbf{pend}(I)$ then r_i is not on the preference list of h_i in I and hence there is no incoming edge to (r_i, h_i) in G_I . Since we are assuming that (r_i, h_i) is not a root of G_I , we conclude that $(r_i, h_i) \in \mathbf{tent}(I) \setminus \mathbf{pend}(I)$. Due to the invariant, (r_i, h_i) is in $\mathbf{rej}(\tau)$ and hence its rejection must be preceded in τ by a proposal of some match (r, h_i) in $\mathbf{prop}(\tau) \subseteq P$ such that r precedes r_i in the preference list of h_i and hence there is an edge of G_i from (r, h_i) to (r_i, h_i) . We let $(r_{i+1}, h_{i+1}) = (r, h_i)$. Next suppose $(r_i, h_i) \in P$. Then, by the invariant we have $(r_i, h_i) \in \mathbf{prop}(\tau)$. If $r_i \notin \mathbf{res tent}(I)$ then (r_i, h_i) is a root of G_I and we are done. Otherwise, the proposal of (r_i, h_i) must be preceded in τ by the rejection of (r_i, h) for some h . We let $(r_{i+1}, h_{i+1}) = (r_i, h)$.

As the construction selects matches appearing in τ in the reversed order, it must eventually end at a root of G_I .

For the converse, suppose there is a directed path p from some root of G_I to (r_0, h_0) . Let τ_p be an event sequence listing the matches in p in the same order and making each match in P a proposal and each match in T a rejection. Extend I by adding h to the preference list of r , for each $(r, h) \in \mathbf{prop}(\tau_p)$. Furthermore, if the starting vertex (r^*, h^*) of p is in T , which implies that $(r^*, h^*) \in \mathbf{pend}(I)$, complete the preference list of h^* so that r^* gets the lowest rank. Let I' be the resulting extension of I . Let σ be a maximal I -feasible sequence. If $(r^*, h^*) \in T$ then, as the quota of each hospital is one, $\sigma + (r^*, h^*)^-$ is I' -feasible. Otherwise, since $(r^*, h^*) \in P$ and $r^* \notin \mathbf{res tent}(I)$, it follows that $\sigma + (r^*, h^*)^+$ is I' -feasible. By a straightforward induction, we may verify that $\sigma + \tau_p$ is I' -feasible. As $\mathbf{rej}(\tau_p) \subseteq \mathbf{tent}(I)$, $\sigma + \tau_p$ is a simple extension of σ and hence I' is a simple extension of I that rejects (r_0, h_0) . \square

The following theorem is immediate from Theorem 2 and Lemma 5.

Theorem 4. *1-FTM-RM is solvable in polynomial time.*

4.3 Hardness of FTM-RM

In this subsection, we show that 2-FTM-RM, and hence FTM-RM, is coNP-complete. The reduction is from SAT through an intermediate problem we call DIGRAPH-FIRING.

Let G be a digraph and $\theta : V(G) \rightarrow \mathbb{N}$ be a *threshold* function which assigns a non-negative integer $\theta(v)$ to each vertex v of G . A θ -firing of G is a subgraph F of G such that, for each $v \in V(F)$, the in-degree of v in F is at least $\theta(v)$ and the out-degree of v in F is at most 1.

k -DIGRAPH-FIRING

Instance: A triple (G, t, θ) , where G is a digraph, t is a vertex of G , and θ is a

threshold function on $V(G)$ such that $\theta(v) \leq k$ for every $v \in V(G)$.

Question: Does G have a θ -firing that contains t ?

Lemma 6. *2-DAG-FIRING is NP-Complete.*

Proof. That 2-DAG-FIRING is in NP is trivial. We show its NP-hardness by a reduction from SAT. Let S be a set of clauses, X the set of variables of S , and C_1, \dots, C_m the enumeration of clauses in S . Using Proposition 3, we assume that each variable in X appears positively in exactly two clauses and negatively in exactly one clause. For each $x \in X$, let i_x^- denote the index of the clause containing x negatively and let i_x^1 and i_x^2 denote the indices of clauses that contain x positively. We construct a DAG G as follows. $V(G)$ contains distinct vertices a_i and b_i for $1 \leq i \leq m$ and five distinct vertices u_x, v_x, l_x^-, l_x^1 , and l_x^2 for each $x \in X$. The edge set is defined by

$$E(G) = \{(a_i, b_i) \mid 1 \leq i \leq m\} \cup \{(b_i, b_{i+1}) \mid 1 \leq i < m\} \cup \bigcup_{x \in X} E_x,$$

where

$$E_x = \{(u_x, l_x^-), (v_x, l_x^-), (u_x, l_x^1), (v_x, l_x^2), (l_x^-, a_{i_x^-}), (l_x^1, a_{i_x^1}), (l_x^2, a_{i_x^2})\}.$$

We set $t = b_m$. The threshold function θ is such that $\theta(v)$ is the in-degree of v except that $\theta(a_i) = 1$ for $1 \leq i \leq m$. Since the only vertices with indegree possibly larger than two are a_i , $1 \leq i \leq m$, we have $\theta(v) \leq 2$ for every $v \in V(G)$.

See Fig. 2 for an example.




Fig. 2. Reduction from SAT to 2-DAG-FIRING

Observe that, for each $x \in X$, a θ -firing cannot contain either l_x^1 or l_x^2 if it contains l_x^- but can contain both l_x^1 and l_x^2 simultaneously if it does not contain l_x^- . Given this property of the “variable gadgets” in G , it is straightforward to see that there is a mutual conversion between a satisfying assignments of S and a θ -firing of G containing t . \square

Theorem 5. *2-FTM-RM is coNP-complete.*

Proof. That 2-FTM-RM is in coNP is trivial. To show that it is coNP-hard, we give a polynomial time reduction from k -DAG-FIRING to the complement of k -FTM-RM, for each positive integer k . As 2-DAG-FIRING is NP-complete by Lemma 6, the theorem follows.

Fix k . Let (G, t, θ) be an instance of k -DAG-FIRING. Without loss of generality, we assume that t is a sink of G . For each $v \in G$, let $N^-(v)$ denote the set of in-neighbors of v in G and let $V_0 = \{v \in V(G) \mid N^-(v) = \emptyset\}$ denote the set of roots of G . We construct an instance $I = (R, H, \{q_h\}_{h \in H}, \{\pi_h\}_{h \in H}, \{\lambda_r\}_{r \in R})$ as follows. For each $v \in V(G)$, we have a mutually distinct resident r_v and we set $R = \{r_v \mid v \in V(G)\}$. For each non-root vertex $v \in V(G) \setminus V_0$, we have a mutually distinct hospital h_v and we set $H = \{h_v \mid v \in V(G) \setminus V_0\}$. For each non-root vertex $v \in V(G) \setminus V_0$, we set $q_{h_v} = \theta(v)$. For each non-root vertex $v \in V(G) \setminus V_0$, the preference list of h_v lists r_u , $u \in N^-(v)$, in the first $|N^-(v)|$ places in an arbitrary order and then lists r_v as its final element. For each root $v \in V_0$, the preference list of r_v is empty (nothing disclosed). For each non-root vertex $v \in V(G) \setminus V_0$, the preference list of r_v consists of a single entry h_v (only the top preference is disclosed). Finally, the match for which we ask the finalizability is (r_t, h_t) . It is straightforward to verify that I is resident-minimal and that $(r_t, h_t) \in \mathbf{tent}(I)$; in fact we have $(r_v, h_v) \in \mathbf{tent}(I)$ for every $v \in V(G) \setminus V_0$. It is also clear that the quota of each hospital in I is k or smaller.

See Figure 3 for an example.

First suppose that G has a θ -firing F that contains t . We show that then (r_t, h_t) is not finalizable in I . Let I' be obtained from I by adding h_v at the end of the preference list of r_u in I , for each $(u, v) \in F$. Let $v_1, \dots, v_n = t$ be a topologically sorted enumeration of $V(F)$. We define I' -feasible sequence σ_i , $0 \leq i \leq n$, inductively as follows. We will maintain the induction hypothesis that σ_i is I' -feasible and $\mathbf{tent}(\sigma_i) = \{(r_{v_j}, h_{v_j}) \mid v_j \in V(G) \setminus V_0 \text{ and } j > i\} \cup \{(r_{v_j}, h_{v_k}) \mid j \leq i \text{ and } (v_i, v_k) \in E(F)\}$. Let σ_0 be an arbitrary maximal I -feasible sequence. Since I -feasibility implies I' -feasibility and $\mathbf{tent}(I) = \{(r_v, h_v) \mid v \in V(G) \setminus V_0\}$, the induction hypothesis holds for the base case. Suppose $i > 0$. If v_i does not have any outgoing edge in F then set $\sigma_i = \sigma_{i-1}$. Suppose v_i has an outgoing edge (v_i, v_k) in F . Because of the topological ordering, we have $k > i$. If $v_i \in V_0$, then set $\sigma_i = \sigma_{i-1} + (r_{v_i}, h_{v_k})^+$. Since $r_{v_i} \notin \mathbf{restent}(I)$ and h_{v_k} is ranked top in the preference list of r_{v_i} in I' , σ_i is I' -feasible. The induction hypothesis is maintained since we have $(v_i, v_k) \in E(F)$ and $(r_{v_i}, h_{v_k}) \in \mathbf{tent}(\sigma_i)$. On the other hand, if $v_i \in V(F) \setminus V_0$ then set $\sigma_i = \sigma_{i-1} + (r_{v_i}, h_{v_i})^- + (r_{v_i}, h_{v_k})^+$. In this case, the in-degree of v_i in F is at least $\theta(v_i) = q_{h_{v_i}}$. For each in-neighbor v_j of v_i in F , its index $j < i$ and, by the induction hypothesis, we have $(r_{v_j}, h_{v_i}) \in$




Fig. 3. Reduction from 3-DAG-FIRING to the complement of 3-FTM-RM

$\mathbf{tent}(\sigma_{i-1})$. Therefore, the preference list of h_{v_i} in I' has at least $q_{h_{v_i}}$ residents in $\mathbf{res}_h \mathbf{tent}(\sigma_{i-1})$ that precede r_{v_i} and therefore $\sigma'_i = \sigma_{i-1} + (r_{v_i}, h_{v_i})^-$ is I' -feasible. Moreover, since r_{v_i} ranks h_{v_k} immediately after h_{v_i} , $\sigma_i = \sigma'_i + (r_{v_i}, h_{v_k})^+$ is I' -feasible. We also have $\mathbf{tent}(\sigma_i) = \mathbf{tent}(\sigma_{i-1}) \cup \{(r_{v_i}, h_{v_k})\} \setminus \{(r_{v_i}, h_{v_i})\}$ and therefore the induction hypothesis is maintained. This construction leads to a I' -feasible sequence σ_n that rejects (r_t, h_t) . Therefore, (r_t, h_t) is not finalizable in I .

For the converse, suppose (r_t, h_t) is not finalizable in I . Since I is resident-minimal, by Theorem 2, there is some simple extension I' of I that rejects (r_t, h_t) . We assume without loss of generality that I' is resident-minimal: take an appropriate truncation if not. Let $U = V_0 \cup \{v \in V(G) \setminus V_0 \mid (r_v, h_v) \in \mathbf{rej}(\tau)\}$ and let F be the subgraph of G induced by U . We show that F is a θ -firing of G . We first show that the out-degree of each vertex in F is at most 1. Let (u, v) be an arbitrary edge of F . By the definition of G , r_u precedes r_v in the preference list of h_v in I . Since $v \in U$, (r_v, h_v) is rejected by τ , which implies that τ contains the proposal of (r, h_v) for every r that precedes r_v in the preference list of h_v (recall that there are exactly q_{h_v} such residents r), including r_u . Therefore, the extension of the preference list of r_u from I to I' is by h_v . This shows that, for each u , the vertex v such that (u, v) is an edge of F is unique if one exists: the out-degree of each vertex in F is at most 1.

We next show that the in-degree of each vertex v is at least $\theta(v)$. If $v \in V_0$, this is obvious since $\theta(v) = 0$. Suppose $v \in U \setminus V_0$. Then, since τ rejects (r_v, h_v) , this rejection event must be preceded in τ by the proposal of (r_u, h_v) for every

resident r_u in the set of the $q_{h_v} = \theta(v)$ residents that precede r_v in the preference list of h_v . But for each such resident r_u , either u is in V_0 (hence $r_u \notin \mathbf{res\ tent}(I)$) or the rejection of (r_u, h_u) precedes the proposal of (r_u, h_v) in τ . In either case, we have $u \in U$. Therefore, the in-degree of v in F is at least $\theta(v)$. We conclude that F is a θ -firing of G . Since τ rejects (r_t, h_t) , we have $t \in U$. This completes the proof that if (r_t, h_t) is not finalizable then there is a θ -firing of G that contains t . \square

5 A sufficient condition for finalizability

In this section, we introduce a polynomial-time decidable sufficient condition for a match to be finalizable in a given instance. This condition turns out necessary for resident-minimal instances in the stable marriage case, thus providing another proof that 1-FTM-RM is polynomial time solvable (Theorem 4).

Let I be an instance and M a subset of $\mathbf{tent}(I)$. We say that $r \in R$ is *relevant to* $h \in H$ with respect to M if r is matched in M either to h or to no hospital. We say that a match (r, h) in M is *endangered in M with respect to I* if it satisfies the following condition: if $(r, h) \in \mathbf{tent}(I) \setminus \mathbf{pend}(I)$ then the preference list of h in I contains q_h or more residents before h that are relevant to h with respect to M ; if $(r, h) \in \mathbf{pend}(I)$ then the number of residents relevant to h with respect to M is $q_h + 1$ or greater. We denote by $\mathbf{dang}_I(M)$ the set of endangered matches with respect to M in I . We say that the set M is *safe* with respect to I if $\mathbf{dang}_I(M) = \emptyset$. Observe that \mathbf{dang}_I is monotone decreasing in the following sense: if $M \subseteq M' \subseteq \mathbf{tent}(I)$ and $(r, h) \in M \setminus \mathbf{dang}_I(M)$ then $(r, h) \notin \mathbf{dang}_I(M')$.

See Table 1 for an example.

Proposition 8. *Let I be an instance and suppose $M \subseteq \mathbf{tent}(I)$ is safe with respect to I . Then, every match in M is finalizable in I .*

Proof. Let σ be an arbitrary feasible extension of the maximal I -feasible sequence. Let τ be the maximal prefix of σ that does not contain the rejection of any member of M . Since $M \subseteq \mathbf{tent}(I)$, τ is an extension of the maximal I -feasible sequence. Since $M \subseteq \mathbf{tent}(I)$ and τ does not reject any match in M , we have $M \subseteq \mathbf{tent}(\tau)$ and hence $\mathbf{dang}_I(\mathbf{tent}(\tau)) \subseteq \mathbf{dang}_I(M) = \emptyset$ by the monotonicity of \mathbf{dang}_I observed above and the assumption that M is safe. This means that there is no match (r, h) in M such that $\tau + (r, h)^-$ is feasible, since if there is such a match then it would be endangered in $\mathbf{tent}(\tau)$ with respect to I . Therefore, τ must be equal to σ and therefore there is no extension of I that rejects any match in M . \square

Theorem 6. *Let I be an instance. Then, the maximal safe set with respect to I is unique and can be identified in polynomial time.*

Proof. Let $M_0 = \mathbf{tent}(I)$ and $M_i = M_{i-1} \setminus \mathbf{dang}_I(M_{i-1})$ for $i > 0$. Let m be the smallest i such that $M_i = M_{i+1}$. Since $\mathbf{dang}_I(M_m) = \emptyset$, M_m is safe. To

Table 1. An example of a safe set

Preference list π_h of each $h \in H$

Matches in **tent**(I) are in bold face

Matches in M , a safe set, are parenthesized

h	1st	2nd	3rd	4th	5th	6th	7th	8th	9th
X	(a)	i	e	(c)	b	f	d	-	-
Y	i	(g)	a	(b)	d	e	c	-	-
Z	e	b	g	a	(i)	d	-	-	-

Preference list of relevant residents for each h

Each member of M is within the quota of 3 in these lists

h	1st	2nd	3rd	4th	5th	6th	7th	8th	9th
X	(a)		e	(c)		f	d	-	-
Y		(g)		(b)	d	e	c	-	-
Z	e				(i)	d	-	-	-

Preference list of each $r \in R$; matches in **tent**(I)

are in bold face; rejected matches are in braces;

r	1st	2nd	3rd
a	X	-	-
b	Y	X	-
c	{Y}	X	-
d	{X}	-	-
e	Y	X	-
f	{Y}	X	-
g	{X}	Y	-
h	{Y}	-	-
i	Z	-	-

show its maximality, let M be an arbitrary subset of **tent**(I) that is safe with respect to I . We show by induction on i that $M \subseteq M_i$. The base case $i = 0$ is trivial. Suppose $i > 0$. By the induction hypothesis, $M \subseteq M_{i-1}$. Let (r, h) be a match in **dang** $_I(M_{i-1})$. We cannot have $(r, h) \in M$, since if we did then, by the monotonicity of **dang** $_I$, (r, h) would be endangered in M , a contradiction to the assumption that M is safe. Therefore, **dang** $_I(M_{i-1}) \cap M = \emptyset$ and hence $M \subseteq M_i$ holds. Therefore we have $M \subseteq M_m$ and hence M_m is the unique maximal safe set. \square

In the stable marriage resident-minimal case, the above sufficient condition for finalizability turns out necessary as well.

Theorem 7. *Let I be a resident-minimal instance in the stable marriage case. Then, each $(r, h) \in \mathbf{tent}(I)$ is finalizable only if it is in the maximal safe set with respect to I .*

Proof. Fix I and let M_i , $0 \leq i \leq m$, be as defined in the proof of Theorem 6. In particular, M_m is the maximal safe set with respect to I . Fix an arbitrary match $(r, h) \in \mathbf{tent}(I) \setminus M_m$. We show that (r, h) is not finalizable.

Let i be the smallest integer such that $(r, h) \notin M_i$. Since $(r, h) \in M_0 = \mathbf{tent}(I)$, we have $i > 0$. We construct a sequence of matches $(r, h) = (r_i, h_i)$, $(r_{i-1}, h_{i-1}), \dots, (r_0, h_0)$ such that $(r_j, h_j) \in \mathbf{dang}_I(M_{j-1})$ for each j , $1 \leq j \leq i$.

Let $0 < j < i$ and suppose $(r_k, h_k) \in \mathbf{dang}_I(M_{k-1})$ for $i \geq k > j$ has been determined. As $(r_{j+1}, h_{j+1}) \in \mathbf{dang}_I(M_j)$, we have some resident, say r_j , that is relevant to h_{j+1} with respect to M_j and precedes r_{j+1} in the preference list of h_{j+1} in I . Observe here that we cannot have $(r_{j+1}, h_{j+1}) \in \mathbf{pend}(I)$ since if we had then (r_{j+1}, h_{j+1}) would be in **dang** $_I(M_0)$ and hence not in

$\mathbf{dang}_I(M_j) \subseteq M_j$ as $j \geq 1$. Now, r_j is relevant to h_{j+1} with respect to M_j but not with respect to M_{j-1} since (r_{j+1}, h_{j+1}) is not in $\mathbf{dang}_I(M_{j-1})$. Therefore, there is some hospital, say h_j , distinct from h_{j+1} such that $(r_j, h_j) \in M_{j-1} \setminus M_j = \mathbf{dang}_I(M_{j-1})$. Thus, we have determined $(r_j, h_j) \in \mathbf{dang}_I(M_{j-1})$ for the current j and, inductively, for each j , $i \geq j \geq 1$.

As observed above, we have $(r_j, h_j) \in \mathbf{tent}(I) \setminus \mathbf{pend}(I)$ for $i \geq j \geq 1$ and, since I is a stable marriage instance, h_j for $i \geq j \geq 1$ are pairwise distinct. More straightforwardly, r_j for $i \geq j \geq 1$ are pairwise distinct as there is at most one match in $\mathbf{tent}(I)$ involving a particular resident.

Let σ be a maximal I -feasible sequence. We now construct successive extensions I_1, \dots, I_m of I and successive extensions $\sigma_1, \dots, \sigma_m$ of σ , based on the sequence of matches constructed above. Since $(r_1, h_1) \in \mathbf{dang}_I(M_0) = \mathbf{dang}_I(\mathbf{tent}(I))$, either $(r_1, h_1) \in \mathbf{pend}(I)$ or $(r_1, h_1) \in \mathbf{tent}(I) \setminus \mathbf{pend}(I)$. If $(r_1, h_1) \in \mathbf{pend}(I)$ then let I_1 be obtained from I by completing the preference list of h_1 so that r_1 is ranked lowest and let $\sigma_1 = \sigma + (r_1, h_1)^-$. Suppose otherwise that $(r_1, h_1) \in \mathbf{tent}(I) \setminus \mathbf{pend}(I)$ then, since (r_1, h_1) is endangered in $\mathbf{tent}(I)$ with respect to I , there must be some resident, say r_0 , that precedes r_1 in the preference list of h_1 in I and is relevant to h_1 with respect to $\mathbf{tent}(I)$. The latter condition implies that $(r_0, h_1) \notin \mathbf{tent}(I)$. Moreover, (r_0, h_1) is not in $\mathbf{prop}(I)$ since if it were then it would be impossible for (r_1, h_1) to be in $\mathbf{tent}(I)$. Since I is resident-minimal, it follows that h_1 is not in the preference list of r_0 . We let I_1 be obtained from I by appending h_1 to the preference list of r_0 and let $\sigma_1 = \sigma + (r_0, h_1)^+ + (r_1, h_1)^-$. In either case, σ_1 is I_1 -feasible. In general, we maintain the invariant that $(r_j, h_j) \in \mathbf{rej}(\sigma_j)$ and σ_j is I_j -feasible for $1 \leq j \leq m$. Suppose $j > 1$ and I_{j-1} together with σ_{j-1} has been defined. Let I_j be obtained from I_{j-1} by appending h_j to the preference list of r_{j-1} and let $\sigma_j = \sigma_{j-1} + (r_{j-1}, h_j)^+ + (r_j, h_j)^-$. As $(r_{j-1}, h_{j-1}) \in \mathbf{rej}(\sigma_{j-1})$, $\sigma_{j-1} + (r_{j-1}, h_j)^+$ is I_j -feasible. Moreover, since r_{j-1} precedes r_j in the preference list of h_j by construction, σ_j is I_j -feasible. We conclude that $(r, h) = (r_m, h_m)$ is not finalizable in I since the extension I_m of I rejects (r_m, h_m) . \square

It follows as a corollary to Proposition 8, Theorem 6, and Theorem 7 that 1-FTM-RM is polynomial time solvable giving another proof of Theorem 4.

6 Student-supervisor assignment: simulations

In this section, we present some simulation results on the student-supervisor assignment procedure mentioned in the introduction. The purpose is to demonstrate that there are realistic markets in which multi-round matching procedures based on FTM can be effective. Since the statistics from the real market are not publicly disclosable, we resort to simulations.

Real market

We first describe the real student-supervisor market in the author's department. Every student in the final year of undergraduate study takes a full year project

course as a part of the requirement for graduation. Every faculty member in the department supervises a project course. Supervisors have quotas as even as possible that sum up to the total number of students. The assignment procedure takes place in the following steps.

1. Students visit supervisors' labs to see the research activity there and get interviews if interested.
2. Each supervisor submits a rank list of students to the central system. This list must be complete. The grade point information is provided to the supervisors. Typically, supervisors make their rank list based on the grade points and the score from the interviews.
3. The rank lists of the supervisors are not public but are partially disclosed in the following manner: if the rank of student s in the list of supervisor p is within the quota of p , then s is notified of this fact.
4. The first round of matching: each student submits a rank list of length up to 3 and the deferred acceptance algorithm is executed. Among the resulting tentative matches, those found finalizable by the sufficient condition in Section 5 are finalized. Both the student and the supervisor of each finalized match are notified.
5. Each student s without a finalized match is informed of the following: (1) the list of unfilled supervisors (supervisors for which the number of finalized matches is strictly smaller than their quota) and (2) the list of unfilled supervisors p such that the rank of s in the rank list of p , after removing students who are finalized to supervisors other than p , is within the quota of p . Note that, in the circumstances in (2), s must be matched to p in any stable matching provided that s ranks p the highest among all supervisors except those that rejected s in the first round.
6. The second round of matching: each student without a finalized match submit a complete rank list of unfilled supervisors. This rank list must be consistent with the rank list in the first round, in that they agree on the ordering of common entries. Then the deferred acceptance algorithm is executed to complete the assignment.

The final outcome of the two rounds of matching is stable, assuming that the rank lists of students in both rounds are consistent with their true preferences. This assumption might be disputable because of the partial disclosures of the supervisors' rank lists before each round, described above. There is no strategic reason for students to change their preference orders but there may be psychological factors. Though these disclosures are introduced for good reasons, we do not include this ingredient in our simulation, partly because we want to avoid such disputes and partly because it is difficult to model the influence of such disclosures on the preferences of students.

Model

We have a set S of students and a set P of supervisors. To model the diversity of interests of students and of attractiveness of the supervisors, we have a set T

of topics. Besides $|S|$, $|P|$, $|T|$ we have parameters k , σ_1 , σ_2 , and σ_3 to be used below.

We have the following random variables, which are mutually independent except for the relationships explicitly described. Each $s \in S$ has a score g_s which has a normal distribution with mean 0.5 and standard deviation σ_1 . Each student $s \in S$ has an interest value $i_{s,t}$ on each topic $t \in T$ and each supervisor p has attractiveness $a_{p,t}$ on each topic $t \in T$. For each student s , the total interest $\sum_{t \in T} i_{s,t}$ is fixed to 1 and the relative magnitude of $i_{s,t}$, $t \in T$, is proportional to a random variable with mean 0.5 and standard deviation σ_3 . For each supervisor p , the total attractiveness $\sum_{t \in T} a_{p,t}$ has a normal distribution with mean 1 and standard deviation σ_2 , truncated to fit in the interval $[0.5, 1.5]$. Given this total attractiveness, the relative magnitude of $a_{p,t}$, $t \in T$, is proportional to a random variable with mean 0.5 and standard deviation σ_3 .

Based on these random variables, the rank lists of students and supervisors are determined as follows. For each pair of student s and supervisor p , let $\mathbf{attraction}(s, p) = \sum_{t \in T} i_{s,t} a_{p,t}$ denote the attraction between s and p , which is the inner product between the interest vector i_s of s and the attractiveness vector a_p of p .

Each student s uses $\mathbf{attraction}(s, p)$ as a score to rank p in the list. The rank lists of supervisors are based on the grade scores of students and the results of interviews described as follows. Each student s has interviews with top k supervisors in the rank list of s . The score supervisor p uses to rank student s is the grade score g_s if p does not interview s . If p does interview s , then the score is modified to reflect the chemistry catalyzed by the interview. We use a simplest model that the score in this case is $g_s + \mathbf{attraction}(s, p)$.

Procedure

The procedure has a parameter r , a positive integer. In the first round, the top r of the rank list of each student are submitted and the deferred acceptance algorithm is executed on this truncated instance. Among the resulting tentative matches, those found finalizable by the sufficient condition in Section 5 are finalized. We count the number of tentative matches, the number of finalized matches, and the number of supervisors that are completely filled by the finalized matches in the first round.

The second round could be executed using the complete rank lists but, in this study, we are not interested in the final outcomes.

Simulation results

In our simulation, we set $|S| = 100$ and $|P| = 10$, round numbers which are close to the real numbers in the author's department. We also fix the following parameters: $|T| = 4$, $k = 5$, $r = 3$, and $\sigma_1 = \sigma_2 = 0.1$. We try several values of parameter σ_3 , which controls the degree of diversity of the interests of the students and of the attractiveness of supervisors.

We have run the simulation 100 times for each value of σ_3 and recorded the average, minimum, and maximum values of each quantity measured. Table 2 shows the results.

σ_3	tentative matches			finalized matches			finalized/tentative			filled supervisors		
	avg.	min	max	avg.	min	max	avg.	min	max	avg.	min	max
0.1	40.16	30	55	35.72	30	50	0.89	0.71	1.0	3.7	3.0	5.0
0.3	73.57	51	91	67.49	44	88	0.91	0.78	1.0	6.29	4.0	8.0
0.5	82.02	64	94	77.37	58	93	0.94	0.82	1.0	7.08	6.0	9.0
0.7	82.81	69	96	78.68	62	91	0.94	0.81	1.0	7.2	6.0	9.0

Table 2. Simulation results

As σ_3 increases, both the number of tentative matches and the number of finalized matches tend to increase, except that the results for $\sigma_3 = 0.5$ and $\sigma_3 = 0.7$ do not show a significant difference. This tendency is plausible, since the diversity of interests and attractiveness would result in the diversity of preferences.

It might be rather surprising that the ratio of the number of finalized matches over the number of tentative matches is consistently high: on average, it is around or above 90% for all values of σ_3 .

With high diversity of interests and attractiveness ($\sigma_3 = 0.5$, for example), on average, about 77 students out of 100 are finalized after the first round and about 7 supervisors out of 10 are filled with finalized matches.

These numbers are fairly close to those from the real supervisor assignment results in the author’s department. Thus, the savings in the evaluation efforts of the students are enormous. The students finalized in the first round do not need to extend their list beyond the top 3 supervisors and those who are not finalized may concentrate on the small number of unfilled supervisors in the second round.

We do not claim that our model captures the underlying mechanism of the real market well. In particular, the model for the effect of interviews is too simplistic. Nonetheless, the simulations do demonstrate that multi-round stable matching procedure based on FTM can be effective for markets where the preferences of participants are diverse and some prematch process, such as interviews, helps nurturing ties between some pairs through which each side of a pair ranks the other high.

7 Future work

Though the sufficient condition for finalizability given in Section 5 is useful in matching procedures for markets as studied in Section 6, exact determination of finalizability, if can be done with a reasonable amount of computation, would further enhance the merit of the multi-round approach. The characterization of

negative instances for FTM-RM given in Section 4 would be indispensable in developing practical algorithms for exact finalizability.

Applicability of the approach to larger markets such as NRMP is an interesting and challenging topic.

Acknowledgments

The author would like to thank his department colleagues, especially Toshiyuki Tsutsumi, for discussions that were helpful in forming the idea of incremental submissions. He is also thankful to Shuichi Miyazaki for helpful comments on an earlier manuscript and to anonymous reviewers of SAGT 2015 conference whose constructive comments on a submission describing this work in a preliminary stage have lead to the improved results and presentations in the current paper. He also thanks SAGT 2016 reviewers for comments that were helpful in improving the presentation.

References

1. David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
2. Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. W.H.Freeman, New York, 2002.
3. Kazuo Iwama and Shuichi Miyazaki. A survey of the stable marriage problem and its variants. In *Informatics Education and Research for Knowledge-Circulating Society, 2008. ICKS 2008. International Conference on*, pages 131–136. IEEE, 2008.
4. NRMP. Match participation agreement for institutions for all matches opening after june 1, 2014, August 2014.
5. Baharak Rastegari, Anne Condon, Nicole Immorlica, Robert Irving, and Kevin Leyton-Brown. Reasoning about optimal stable matchings under partial information. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 431–448. ACM, 2014.
6. Baharak Rastegari, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the fourteenth ACM conference on Electronic Commerce*, pages 733–750. ACM, 2013.
7. Alvin E Roth and Uriel G Rothblum. Truncation strategies in matching markets-in search of advice for participants. *Econometrica*, 67(1):21–43, 1999.